arXiv:1511.01716v3 [math.OC] 12 Aug 2016

Selective Linearization For Multi-Block Convex Optimization

Yu Du* Xiaodong Lin† Andrzej Ruszczyński‡

November 6, 2015; revised: March 3, 2016; August 12, 2016

Abstract

We consider the problem of minimizing a sum of several convex non-smooth functions. We introduce a new algorithm called the selective linearization method, which iteratively linearizes all but one of the functions and employs simple proximal steps. The algorithm is a form of multiple operator splitting in which the order of processing partial functions is not fixed, but rather determined in the course of calculations. Global convergence is proved and estimates of the convergence rate are derived. Specifically, the number of iterations needed to achieve solution accuracy ε is of order $\mathcal{O}\big(\ln(1/\varepsilon)/\varepsilon\big)$. We also illustrate the operation of the algorithm on structured regularization problems.

Keywords: Nonsmooth optimization, operator splitting, multiple blocks, alternating linearization

1 Introduction

In recent years, we have seen extensive development of the theory and methods for *structured regularization*, one of the most fundamental techniques to address the "big data" challenge. The basic problem is to minimize the following objective function with two components (blocks):

$$\min \Big[F(x) = f_1(x) + f_2(x)\Big], \tag{1}$$

where $f_1(\cdot)$ is the loss function and $f_2(\cdot)$ is a penalty function that imposes structured regularization to the model. Both functions are usually convex, but may be nonsmooth. Many data mining and machine learning problems can be cast within this framework, and efficient methods were proposed for these problems. The first group are the *operator*

*Rutgers University, Department of Management Science and Information Systems, Piscataway, NJ 08854, USA, Email:duyu@rutgers.edu

†Rutgers University, Department of Management Science and Information Systems, Piscataway, NJ 08854, USA, Email:linxd@rci.rutgers.edu

‡Rutgers University, Department of Management Science and Information Systems, Piscataway, NJ 08854, USA, Email:rusz@rutgers.edu

splitting methods originating from [DR56] and [PR55], and later developed and analyzed by [BC11, Com09, EB92, LM79], among others. Their dual versions, known as *Alternating Direction Methods of Multipliers* (ADMM) (see, [GM76, GM75, GLT89]), found many applications in signal processing (see, *e.g.*, [BPC$^+$10, CP11], and the references therein). Sometimes, they are called *split Bregman methods* (see, *e.g.*, [GO09, YX11]).

The *Alternating Linearization Method* (ALIN) of [KRR99] handles problems of form (1) by introducing an additional improvement test to the operator splitting methods, which decides whether the proximal center should be updated or stay unchanged, and which of the operator splitting formulas should be applied at the current iteration. Its convergence mechanism is different than that of the splitting methods; it adapts some ideas of bundle methods of nonsmooth optimization [HUL93, Kiw85, Rus06]. The recent application of ALIN to structured regularization problems in [LPR14] proved very successful, with fast convergence, good accuracy, and scalability to very large dimensions.

Most of existing techniques for structured regularization are designed to handle the two-block problem of form (1).

In this paper, we plan to extend the ALIN framework to optimization problems involving multiple components. Namely, we aim to solve the following problem:

$$\min \Big\{ F(x) = \sum_{i=1}^{N} f_i(x) \Big\}, \tag{2}$$

with convex (possibly nondifferentiable) functions $f_i : \mathbb{R}^n \to \mathbb{R}$, $i = 1, \ldots, N$, where the number of component functions, N, may be arbitrarily large. We only assume that the minimum exists.

In a typical application, $f_1(\cdot)$ may be the loss function, similar to problem (1), while the penalty function is a sum of multiple components. This type of generalization has many practical applications, including low rank matrix completion, compressed sensing, dynamic network analysis, and computer vision.

To the best of the authors' knowledge, no general convergent versions of operator splitting methods for multiple blocks exist. A known way is to introduce N copies $x^1 = x^2 = \cdots = x^N$ of x, and reduce the problem to the two-function case in the space $\mathbb{R}^{nN}$ [CP11]:

$$\min \sum_{i=1}^{N} f_i(x^i) + I(x^1, \ldots, x^N)$$

with $I(\cdot)$ denoting the indicator function of the subspace $x^1 = x^2 = \cdots = x^N$. Similar ideas were used in stochastic programming, under the name of *Progressive Hedging* [RW91]. A method for three blocks with one function being differentiable was theoretically analyzed in [Con13].

Our new algorithm, which we call the *Selective Linearization Method* (SLIN), does not replicate the decision variables. It generates a sequence of points $x^k \in \mathbb{R}^n$ with a monotonic sequence of corresponding function values $\{F(x^k)\}$. At each iteration, it linearizes all but one of the component functions and uses a proximal term penalizing for the distance to the last iterate. In a sense, each step is a backward step of the form

employed in operator splitting. The order of processing the functions is not fixed; the method uses precise criteria for selecting the function to be treated exactly at the current step. It also employs special rules for updating the proximal center. These two rules differ our approach from the simultaneously proposed incremental proximal method of [Ber15], which applies to smooth functions only, and achieves linear convergence rate in this case.

The algorithm is a multi-block extension of the Alternating Linearization method for solving two-block nonsmooth optimization problems. Global convergence and convergence rate of the new algorithm are proved. Specifically, the new algorithm is proven to require at most $\mathcal{O}\big(\ln(1/\varepsilon)/\varepsilon\big)$ iterations to achieve solution accuracy ε.

In section 2, we present the method and prove its global convergence. The convergence rate is derived in section 4. Finally, in section 5, we illustrate its operation on structured regularized regression problems involving many blocks.

2 The Method

Our method derives from two fundamental ideas of convex optimization: the *Moreau–Yosida regularization* of $F(\cdot)$,

$$F_D(y) = \min\Big\{F(x) + \frac{1}{2}\big\|x-y\big\|_D^2\Big\}, \tag{3}$$

and the *proximal step* for (2),

$$\operatorname{prox}_F(y) = \arg\min\Big\{F(x) + \frac{1}{2}\big\|x-y\big\|_D^2\Big\}. \tag{4}$$

In the formulas above, the norm $\|x\|_D = \big(\langle x, Dx\rangle\big)^{1/2}$ with a positive definite matrix D. In applications, we shall use a diagonal D, which leads to major computational simplifications. The *proximal point method* carries out the iteration $x^{k+1} = \operatorname{prox}_F(x^k)$, $k = 1, 2, \ldots$ and is known to converge to a minimum of $F(\cdot)$, if a minimum exists [Roc76].

The main idea of our method is to replace problem (2) with a sequence of approximate problems of the following form:

$$\min_x f_{j_k}(x) + \sum_{i\neq j_k} \tilde{f}_i^k(x) + \frac{1}{2}\big\|x-x^k\big\|_D^2. \tag{5}$$

Here $k = 1, 2, \ldots$ is the iteration number, x^k is the current best approximation to the solution, $j_k \in \{1, \ldots, N\}$ is an index selected at iteration k, and $\tilde{f}_i^k$ are *affine minorants* of the functions f_i, $i \in \{1, \ldots, N\} \setminus \{j_k\}$. These minorants are constructed as follows:

$$\tilde{f}_i^k(x) = f_i(z_i^k) + \langle g_i^k, x - z_i^k\rangle,$$

with some points $z_i^k \in \mathbb{R}^n$ and specially selected subgradients $g_i^k \in \partial f_i(z_i^k)$. Thus, problem (5) differs from the proximal point problem in (3) by the fact that only one of

the functions $f_i(\cdot)$ is treated exactly, while the other functions are replaced by affine approximations.

The key elements of the method are the selection of the index j_k, the way the affine approximations are constructed, and the update rule for the proximal center x^k. In formula (5) and in the algorithm description below we write simply $i \neq j_k$ for $i \in \{1, \ldots, N\} \setminus \{j_k\}$.

We denote the function approximating $F(x)$ in (5) by

$$\tilde{F}^k(x) = f_{j_k}(x) + \sum_{i \neq j_k} \tilde{f}_i^k(x).$$

Selective Linearization (SLIN) Algorithm

Step 0: Set $k = 1$ and $j_1 \in \{1, \ldots, N\}$, select $x^1 \in \mathbb{R}^n$ and, for all $i \neq j_1$, linearization points $z_i^1 \in \mathbb{R}^n$ where the corresponding subgradients $g_i^1 \in \partial f_i(z_i^1)$ exist. Define $\tilde{f}_i^1(x) = f_i(z_i^1) + \langle g_i^1, x - z_i^1 \rangle$ for $i \neq j_1$. Choose parameters $\beta \in (0, 1)$, and a stopping precision $\varepsilon > 0$.

Step 1: Find the solution $z_{j_k}^k$ of the f_{j_k}-subproblem (5) and define

$$g_{j_k}^k = -\sum_{i \neq j_k} g_i^k - D(z_{j_k}^k - x^k). \tag{6}$$

Step 2: If

$$F(x^k) - \tilde{F}^k(z_{j_k}^k) \leq \varepsilon, \tag{7}$$

then stop. Otherwise, continue.

Step 3: If

$$F(z_{j_k}^k) \leq F(x^k) - \beta\big(F(x^k) - \tilde{F}^k(z_{j_k}^k)\big), \tag{8}$$

then set $x^{k+1} = z_{j_k}^k$ (*descent step*); otherwise set $x^{k+1} = x^k$ (*null step*).

Step 4: Select

$$j_{k+1} = \arg\max_{i \neq j_k} \big\{ f_i(z_{j_k}^k) - \tilde{f}_i^k(z_{j_k}^k) \big\}. \tag{9}$$

For all $i \neq j_{k+1}$, set $z_i^{k+1} = z_i^k$ and $g_i^{k+1} = g_i^k$ (so that $\tilde{f}_i^{k+1}(\cdot) \equiv \tilde{f}_i^k(\cdot)$). Increase k by 1 and go to Step 1.

Few comments are in order. Since the point $z_{j_k}^k$ is a solution of the subproblem (5), the vector $g_{j_k}^k$ calculated in (6) is indeed a subgradient of f_{j_k} at $z_{j_k}^k$; in fact, it is exactly the subgradient that features in the optimality condition for (5) at $z_{j_k}^k$. Therefore, at all iterations, the functions $\tilde{f}_i^k(\cdot)$ are minorants of the functions $f_i(\cdot)$. This in turn implies that $\tilde{F}^k(\cdot)$ is a lower approximation of $F(\cdot)$. Consequently, $F(x^k) - \tilde{F}^k(z_{j_k}^k) \geq 0$ in (7), with $F(x^k) = \tilde{F}^k(z_{j_k}^k)$ equivalent to x^k being the minimizer of $F(\cdot)$.

In practical implementation of the algorithm, the points z_i^k need not be stored. It is sufficient to memorize $\alpha_i^k = f_i(z_i^k) - \langle g_i^k, z_i^k \rangle$ and the subgradients g_i^k. At Step 4,

we then set $\alpha_i^{k+1} = \alpha_i^k$ and $g_i^{k+1} = g_i^k$ for all $i \in \{1, \ldots, N\} \setminus \{j_k, j_{k+1}\}$, while $\alpha_{j_k}^{k+1} = f_i(z_i^k) - \langle g_i^k, z_i^k \rangle$. For j_{k+1} these data are not needed, because the function $f_{j_{k+1}}(\cdot)$ will not be linearized at the next iteration.

In some cases, the storage of the subgradients g_i^k may be substantially simplified.

Example 2.1. Suppose

$$F(x) = \sum_{i=1}^{N} \varphi_i(a_i^T x),$$

with convex functions $\varphi_i : \mathbb{R} \to \mathbb{R}$ and $a_i \in \mathbb{R}^n$, $i-1, \ldots, n$. Then every subgradient of $f_i(x) = \varphi_i(a_i^T x)$ has the form $g_i^k = \sigma_i^k a_i$, with $\sigma_i^k \in \partial\varphi_i(a_i^T z_i^k)$. The scalars σ_i^k are sufficient for recovering the subgradients, because the vectors a_i are part of the problem data.

3 Global convergence

We assume that $\varepsilon = 0$ in Step 2. To prove convergence of the algorithm, we consider two cases: with finitely or infinitely many descent steps.

We first address the finite case and show that the proximal center updated in the last descent step must be an optimal solution to problem (2). To this end, we prove that if a null step is made at iteration k, then the optimal objective function values of consecutive subproblems are increasing and the gap is bounded below by a value determined by

$$v_k = F(x^k) - \tilde{F}^k(z_{j_k}^k). \tag{10}$$

We shall also use this result in the proof of convergence rate.

We denote the optimal objective function value of subproblem (5) at iteration k by

$$\eta^k = \min_x f_{j_k}(x) + \sum_{i \neq j_k} \tilde{f}_i^k(x) + \frac{1}{2}\|x - x^k\|_D^2.$$

Lemma 3.1. *If a null step is made at iteration k, then*

$$\eta^{k+1} \geq \eta^k + \frac{1-\beta}{2(N-1)} \bar{\mu}_k v_k, \tag{11}$$

where

$$\bar{\mu}_k = \min\Big\{1, \frac{(1-\beta)v_k}{(N-1)\|s_{j_{k+1}}^k - g_{j_{k+1}}^k\|_{D^{-1}}^2}\Big\}, \tag{12}$$

with an arbitrary $s_{j_{k+1}}^k \in \partial f_{j_{k+1}}(z_{j_k}^k)$.

Proof. The change from the f_{j_k}-subproblem to the $f_{j_{k+1}}$-subproblem can be viewed as two steps: first is the change of $f_{j_k}(\cdot)$ to $\tilde{f}_{j_k}^k(\cdot)$, followed by the change of $\tilde{f}_{j_{k+1}}^k(\cdot)$ to $f_{j_{k+1}}(\cdot)$. By the selection of the subgradient (6) and the resulting construction of $\tilde{f}_{j_k}^k(\cdot)$,

the first operation does not change the solution and the optimal value of the subproblem. Thus the optimal value of (5) satisfies the following equation:

$$\eta^k = \min_x \sum_{i=1}^{N} \tilde{f}_i^k(x) + \frac{1}{2}\|x - x^k\|_D^2. \tag{13}$$

Since $x^{k+1} = x^k$ at a null step, and $f_{j_{k+1}} \geq \tilde{f}_{j_{k+1}}^k$, the second operation can only increase the optimal value of the last problem. Therefore, $\eta^{k+1} \geq \eta^k$.

Consider the family of relaxations of the $f_{j_{k+1}}$-subproblem at iteration $k+1$:

$$\begin{aligned}\hat{Q}_k(\mu) = \min_x \Bigg\{ & \sum_{i \neq j_{k+1}} \tilde{f}_i^{k+1}(x) + (1-\mu)\left(f_{j_{k+1}}(z_{j_{k+1}}^k) + \langle g_{j_{k+1}}^k, x - z_{j_{k+1}}^k\rangle\right) \\ & + \mu\left(f_{j_{k+1}}(z_{j_k}^k) + \langle s_{j_{k+1}}^k, x - z_{j_k}^k\rangle\right) + \frac{1}{2}\|x - x^k\|_D^2\Bigg\},\end{aligned} \tag{14}$$

with parameter $\mu \in [0,1]$. In the above relaxation, the function $f_{j_{k+1}}(\cdot)$ is replaced by a convex combination of its two affine minorants: one at the point $z_{j_{k+1}}^k$, which is $\tilde{f}_{j_{k+1}}^k(\cdot)$ used at iteration k, and the other one at the kth trial point $z_{j_k}^k$, with an arbitrary subgradient $s_{j_{k+1}}^k$. Due to (13), the value of (14) with $\mu = 0$ coincides with η^k. Therefore, the difference between η^{k+1} and η^k can be estimated from below by the increase in the optimal value $\hat{Q}_k(\mu)$ of (14) when μ moves away from zero. That is,

$$\eta^{k+1} - \eta^k \geq \max_{\mu \in [0,1]} \hat{Q}_k(\mu) - \hat{Q}_k(0). \tag{15}$$

Define $\delta_k = F(z_{j_k}^k) - \tilde{F}^k(z_{j_k}^k)$. Note that $\delta_k \geq 0$, since $f_i \geq \tilde{f}_i^k$ for $i \neq j_k$. We also define $\mu_k = \min\left\{1, \frac{\delta_k}{(N-1)\|s_{j_{k+1}}^k - g_{j_{k+1}}^k\|_{D^{-1}}^2}\right\}$, so $\mu_k \in [0,1]$. By direct calculation, and with the use of (6), the solution of (14) has the form

$$\hat{x}(\mu) = x^k - D^{-1}\Big[\sum_{i=1}^{N} g_i^k + \mu\big(s_{j_{k+1}}^k - g_{j_{k+1}}^k\big)\Big] = z_{j_k}^k + \mu D^{-1}\big(s_{j_{k+1}}^k - g_{j_{k+1}}^k\big).$$

Using the definitions following (14) and the fact that $\hat{x}(0) = z_{j_k}^k$, the derivative of $\hat{Q}_k$ can be expressed as follows:

$$\begin{aligned}\hat{Q}_k'(\mu) &= \langle s_{j_{k+1}}^k - g_{j_{k+1}}^k, \hat{x}(\mu)\rangle \\ &\quad + \left(f_{j_{k+1}}(z_{j_k}^k) - \langle s_{j_{k+1}}^k, z_{j_k}^k\rangle\right) - \left(f_{j_{k+1}}(z_{j_{k+1}}^k) - \langle g_{j_{k+1}}^k, z_{j_{k+1}}^k\rangle\right) \\ &= \langle s_{j_{k+1}}^k - g_{j_{k+1}}^k, \hat{x}(\mu) - z_{j_k}^k\rangle \\ &\quad + f_{j_{k+1}}(z_{j_k}^k) - \left(f_{j_{k+1}}(z_{j_{k+1}}^k) + \langle g_{j_{k+1}}^k, z_{j_k}^k - z_{j_{k+1}}^k\rangle\right) \\ &\geq \langle s_{j_{k+1}}^k - g_{j_{k+1}}^k, \hat{x}(\mu) - z_{j_k}^k\rangle + \frac{F(z_{j_k}^k) - \tilde{F}^k(z_{j_k}^k)}{N-1} \\ &= -\mu\|s_{j_{k+1}}^k - g_{j_{k+1}}^k\|_{D^{-1}}^2 + \frac{\delta_k}{N-1}.\end{aligned} \tag{16}$$

In the inequality above, we used the definition (9) of j_{k+1} and the fact that the maximum of the differences $f_j(z^k_{j_k}) - \tilde f^k_j(z^k_{j_k})$ over $j \neq j_k$ is larger than their average. Thus

(17)
$$\hat Q_k(\mu_k) - \hat Q_k(0) = \int_0^{\mu_k} \hat Q_k'(\mu)\, d\mu \geq \mu_k \left(\frac{\delta_k}{N-1} - \frac{1}{2}\mu_k \big\| s^k_{j_{k+1}} - g^k_{j_{k+1}} \big\|^2_{D^{-1}} \right).$$

Substitution of the definition of μ_k yields

$$\eta^{k+1} \geq \eta^k + \frac{\mu_k \delta_k}{2(N-1)}. \tag{18}$$

If a null step is made at iteration k, then the update step rule (8) is violated. Thus, $\delta_k = F(z^k_{j_k}) - \tilde F^k(z^k_{j_k}) > (1-\beta) v_k$. Plugging this lower bound on δ_k into (18) and using the definition of $\bar\mu_k$, we obtain the postulated bound (11).

Finally, we remark that $s^k_{j_{k+1}} \neq g^k_{j_{k+1}}$, because $f_{j_{k+1}}(z^k_{j_k}) > \tilde f^k_{j_{k+1}}(z^k_{j_k})$. □

We also need to estimate the size of the steps made by the method.

Lemma 3.2. *At every iteration k,*

$$\frac{1}{2}\big\| z^k_{j_k} - \operatorname{prox}_F(x^k) \big\|^2_D \leq F_D(x^k) - \eta^k. \tag{19}$$

Proof. Since $F(\cdot) \geq \tilde F^k(\cdot)$ and $z^k_{j_k}$ is a solution of the strongly convex problem (5), we have

$$\begin{aligned} F_D(x^k) &= F\big(\operatorname{prox}_F(x^k)\big) + \frac{1}{2}\big\| \operatorname{prox}_F(x^k) - x^k \big\|^2_D \\ &\geq \tilde F^k\big(\operatorname{prox}_F(x^k)\big) + \frac{1}{2}\big\| \operatorname{prox}_F(x^k) - x^k \big\|^2_D \\ &\geq \tilde F^k(z^k_{j_k}) + \frac{1}{2}\big\| z^k_{j_k} - x^k \big\|^2_D + \frac{1}{2}\big\| z^k_{j_k} - \operatorname{prox}_F(x^k) \big\|^2_D \\ &= \eta^k + \frac{1}{2}\big\| z^k_{j_k} - \operatorname{prox}_F(x^k) \big\|^2_D. \end{aligned} \tag{20}$$

Rearranging, we obtain (19). □

We are now ready to prove optimality in the case of finitely many descent steps.

Theorem 3.3. *Suppose $\varepsilon = 0$, the set $\mathcal{K} = \{1\} \cup \{k > 1 : x^k \neq x^{k-1}\}$ is finite and $\inf F > -\infty$. Let $m \in \mathcal{K}$ be the largest index such that $x^m \neq x^{m-1}$. Then $x^m \in \operatorname{Argmin} F$.*

Proof. We argue by contradiction. Suppose $x^m \notin \operatorname{Argmin} F$. If $\varepsilon = 0$ the method cannot stop, because $\tilde F^k(z^k_{j_k}) \leq F(\operatorname{prox}_F(x^m)) < F(x^m)$, for all $k \geq m$. Therefore, null steps are made at all iterations $k \geq m$, with $x^k = x^m$. By Lemma 3.1, the sequence $\{\eta^k\}$ is nondecreasing and bounded above by $F(x^m)$. Hence $\eta^{k+1} - \eta^k \to 0$. The right hand side of estimate (19) with $x^k = x^m$ for $k \geq m$, owing to the monotonicity of $\{\eta^k\}$, is nonincreasing, and thus the sequence $\{z^k_{j_k}\}$ is bounded. Since the subgradients of a convex function are locally bounded, the differences $\big\| s^k_{j_{k+1}} - g^k_{j_{k+1}} \big\|_{D^{-1}}$ appearing in

the definition of $\bar{\mu}_k$ in Lemma 3.1 are bounded from above. Therefore, $v_k \to 0$. As $F(x^m) \geq \eta^k \geq F(x^m) - v_k$, we have $\eta^k \uparrow F(x^m)$.

On the other hand, the inequality $\tilde{F}^k(\cdot) \leq F(\cdot)$ implies that $\eta^k \leq F_D(x^m)$ for all $k \geq m$. Since $x^m \notin \operatorname{Argmin} F$, we have $F_D(x^m) < F(x^m)$, which contradicts the convergence of $\{\eta^k\}$ to $F(x^m)$. □

We now address the infinite case. Note that the update test (8) can be expressed as follows:

$$\tilde{F}^k(z^k_{j_k}) \geq -\frac{1}{\beta}F(z^k_{j_k}) + \frac{1-\beta}{\beta}F(x^k). \tag{21}$$

Theorem 3.4. *Suppose* $\operatorname{Argmin} F \neq \emptyset$. *If the set* $\mathcal{K} = \{k : x^{k+1} \neq x^k\}$ *is infinite, then* $\lim_{k\to\infty} x^k = x^*$, *for some* $x^* \in \operatorname{Argmin} F$.

Proof. Consider iteration $k \in \mathcal{K}$ (descent step). From the optimality condition for (5) we obtain

$$0 \in \partial\Big[f_{j_k}(z^k_{j_k}) + \sum_{i\neq j_k} \tilde{f}^k_i(z^k_{j_k})\Big] + D\big(z^k_{j_k} - x^k\big), \tag{22}$$

which yields

$$D(x^k - x^{k+1}) \in \partial\Big[f_{j_k}(z^k_{j_k}) + \sum_{i\neq j_k} \tilde{f}^k_i(z^k_{j_k})\Big]. \tag{23}$$

Then for any point $x^* \in \operatorname{Argmin} F$ we obtain

$$F(x^*) \geq \tilde{F}^k(x^*) \geq \tilde{F}^k(x^{k+1}) + \big\langle D(x^k - x^{k+1}), x^* - x^{k+1}\big\rangle. \tag{24}$$

Hence

$$\begin{aligned}
\big\|x^{k+1} - x^*\big\|^2_D &= \big\|x^k - x^*\big\|^2_D + 2\big\langle D(x^{k+1} - x^k), x^k - x^*\big\rangle + \big\|x^{k+1} - x^k\big\|^2_D \\
&\leq \big\|x^k - x^*\big\|^2_D + 2\big\langle D(x^{k+1} - x^k), x^{k+1} - x^*\big\rangle \\
&\leq \big\|x^k - x^*\big\|^2_D + 2\big(F(x^*) - \tilde{F}^k(x^{k+1})\big).
\end{aligned}$$

Using (21), we can continue this chain of inequalities as follows

$$\begin{aligned}
\big\|x^{k+1} - x^*\big\|^2_D &\leq \big\|x^k - x^*\big\|^2_D + 2\big(F(x^*) - \frac{1}{\beta}F(x^{k+1}) + \frac{1-\beta}{\beta}F(x^k)\big) \\
&= \big\|x^k - x^*\big\|^2_D + 2\big(F(x^*) - F(x^k)\big) + \frac{2}{\beta}\big(F(x^k) - F(x^{k+1})\big).
\end{aligned} \tag{25}$$

Thus, adding up (25) for all $k \in \mathcal{K}$, $k \leq m$, and noting that the null steps do not change the proximal centers, we obtain

$$\big\|x^{m+1} - x^*\big\|^2_D \leq \big\|x^1 - x^*\big\|^2_D + 2\sum_{\substack{k\in\mathcal{K}\\ k\leq m}} \big(F(x^*) - F(x^k)\big) + \frac{2}{\beta}\sum_{\substack{k\in\mathcal{K}\\ k\leq m}} \big(F(x^k) - F(x^{k+1})\big). \tag{26}$$

The term $2\sum_{k\in\mathcal{K},k\le m}\big(F(x^*)-F(x^k)\big)$ is non-positive, and the last term is bounded by $\frac{2}{\beta}\big(F(x^1)-F(x^*)\big)$. Thus, several conclusions follow from inequality (26). First, the sequence $\{x^k\}_{k\in\mathcal{K}}$ is bounded, because their distances to x^* are bounded. Secondly, rewriting (26) as

$$\sum_{k\in\mathcal{K},k\le m}\big(F(x^k)-F(x^*)\big)\le\frac{1}{2}\big(\|x^1-x^*\|_D^2-\|x^{m+1}-x^*\|_D^2\big)+\frac{1}{\beta}\big(F(x^1)-F(x^{m+1})\big),$$

and letting $m\to\infty$ in $\mathcal{K}$, we deduce that

$$\sum_{k\in\mathcal{K}}\big(F(x^k)-F(x^*)\big)\le\frac{1}{2}\|x^1-x^*\|_D^2+\frac{1}{\beta}\big(F(x^1)-F(x^*)\big). \tag{27}$$

Consequently, $F(x^k)\to F(x^*)$ as $k\to\infty$ in $\mathcal{K}$. As the null steps do not change the proximal centers, we also have $F(x^k)\to F(x^*)$, when $k\to\infty$.

To prove that the sequence of proximal centers converges to an optimal solution, note that since the infinite sequence $\{x^k\}_{k\in\mathcal{K}}$ is bounded, it has a convergent subsequence whose limit $\hat{x}$ is a minimizer of F. Without loss of generality, we substitute $\hat{x}$ for x^* in the above derivations, and add (25) for all $k\in\mathcal{K}$ such that $\ell\le k\le m$. For any $1\le\ell\le m$ we obtain the following analog of (26):

$$\begin{aligned}\|x^{m+1}-\hat{x}\|_D^2&\le\|x^\ell-\hat{x}\|_D^2+2\sum_{\substack{k\in\mathcal{K}\\k\le m}}\big(F(\hat{x})-F(x^k)\big)+\frac{2}{\beta}\sum_{\substack{k\in\mathcal{K}\\k\le m}}\big(F(x^k)-F(x^{k+1})\big)\\&\le\|x^\ell-\hat{x}\|_D^2+\frac{2}{\beta}\big(F(x^\ell)-F(\hat{x})\big).\end{aligned}$$

The right hand side of the last inequality can be made arbitrarily small by choosing ℓ from the subsequence converging to $\hat{x}$. Therefore the entire sequence $\{x^k\}_{k\in\mathcal{K}}$ is convergent to $\hat{x}$. □

We finish this section with a number of conclusions, which will be useful in the analysis of the rate of convergence.

Lemma 3.5. *If there is a descent step at iteration k, then*

$$\eta^{k+1}-\eta^k\ge-\|x^{k+1}-x^k\|_D^2\ge\frac{1}{\beta}\big(F(x^{k+1})-F(x^k)\big). \tag{28}$$

Proof. By (6),

$$\sum_{i=1}^N g_i^k+D(x^{k+1}-x^k)=0. \tag{29}$$

The optimal value of (5) at iteration $k+1$ can be then estimated as follows:

$$\begin{aligned}\eta^{k+1} &= \min_x \left\{ f_{j_{k+1}}(x) + \sum_{i \neq j_{k+1}} \tilde{f}_i^{k+1}(x) + \frac{1}{2}\|x - x^{k+1}\|_D^2 \right\} \\ &\geq \min_x \left\{ \sum_{i=1}^N \tilde{f}_i^{k+1}(x) + \frac{1}{2}\|x - x^{k+1}\|_D^2 \right\} \\ &= \sum_{i=1}^N \tilde{f}_i^{k+1}(x^{k+1}) + \min_x \left\{ \Big\langle \sum_{i=1}^N g_i^k, x - x^{k+1} \Big\rangle + \frac{1}{2}\|x - x^{k+1}\|_D^2 \right\} \\ &= \tilde{F}^k(x^{k+1}) + \min_x \left\{ -\big\langle D(x^{k+1} - x^k), x - x^{k+1} \big\rangle + \frac{1}{2}\|x - x^{k+1}\|_D^2 \right\}.\end{aligned}$$

The minimizer on the right hand side is $x = 2x^{k+1} - x^k$, and we conclude that

$$\eta^{k+1} \geq \tilde{F}^k(x^{k+1}) - \frac{1}{2}\|x^{k+1} - x^k\|_D^2 = \eta^k - \|x^{k+1} - x^k\|_D^2,$$

which proves the left inequality in (28). To prove the right inequality, we observe that the test (8) for the descent step is satisfied at iteration k, and thus

$$F(x^k) - F(x^{k+1}) \geq \beta\big(F(x^k) - \tilde{F}^k(x^{k+1})\big) = \beta\big(\tilde{F}^k(x^k) - \tilde{F}^k(x^{k+1})\big).$$

The expression on the right hand side can be calculated with the use of (29), exactly as in the derivations above, which yields

$$F(x^k) - F(x^{k+1}) \geq \beta\|x^{k+1} - x^k\|_D^2.$$

This proves the right inequality in (28). □

We can now summarize convergence properties of the sequences generated by the algorithm.

Corollary 3.6. *Suppose* $\operatorname{Argmin} F \neq \emptyset$ *and* $\varepsilon = 0$. *Then a point* $x^* \in \operatorname{Argmin} F$ *exists, such that:*
(i) $\lim_{k\to\infty} x^k = \lim_{k\to\infty} z_{j_k}^k = x^*$;
(ii) $\lim_{k\to\infty} \eta^k = F(x^*)$.

Proof. The convergence of $\{x^k\}$ to a minimum point x^* has been proved in Theorems 3.3 and 3.4. It remains to verify the convergence properties of $\{z_{j_k}^k\}$ and $\{\eta^k\}$. It follows from Lemmas 3.1 and 3.5 that the sequence $\eta^k - \frac{1}{\beta}F(x^k)$ is nondecreasing. Since $\eta^k \leq F(x^k)$ by construction, this sequence is bounded from above, and thus convergent. Therefore, a limit η^* of $\{\eta^k\}$ exists and $\eta^* \leq F(x^*)$. If the number of descent steps is finite, the equality $\eta^* = F(x^*)$ follows from Theorem 3.3. If the number of descent steps is infinite, inequality (8) at each descent step k yields:

$$F(x^k) - \eta^k \leq F(x^k) - \tilde{F}^k(x^{k+1}) \leq \frac{1}{\beta}\big(F(x^k) - F(x^{k+1})\big).$$

Passing to the limit over descent steps $k \to \infty$ we conclude that $\eta^* \geq F(x^*)$. Consequently, $\eta^* = F(x^*)$ and assertion (ii) is true.

The convergence of the sequence $\{z^k_{j_k}\}$ to x^* follows from inequality (19), because $x^k \to x^*$ and $\eta^k \to F(x^*)$. □

4 Rate of Convergence

Our objective in this section is to estimate the rate of convergence of the method. To this end, we assume that $\varepsilon > 0$ at Step 2 (inequality (7)) and we estimate the number of iterations needed to achieve this accuracy. We also make an additional assumption about the growth rate of the function $F(\cdot)$.

Assumption 4.1. *The function $F(\cdot)$ has a unique minimum point x^* and a constant $\alpha > 0$ exists, such that*

$$F(x) - F(x^*) \geq \alpha \big\|x - x^*\big\|^2_D,$$

for all $x \in \mathbb{R}^n$.

Assumption 4.1 has a number of implications on the properties of the method. First, we recall from [Rus06, Lem. 7.12] the following estimate of the Moreau–Yosida regularization.

Lemma 4.2. *For any point $x \in \mathbb{R}^n$, we have*

$$F_D(x) \leq F(x) - \big\|x - x^*\big\|^2_D \, \varphi\bigg(\frac{F(x) - F(x^*)}{\big\|x - x^*\big\|^2_D}\bigg), \tag{30}$$

where

$$\varphi(t) = \begin{cases} t^2 & \text{if } t \in [0,1], \\ -1 + 2t & \text{if } t \geq 1. \end{cases}$$

Proof. See [Rus06, Lem. 7.12]. □

Lemma 4.3. *Suppose Assumption 4.1 is satisfied. Then the stopping test* (7) *implies that*

$$F(x^k) - F(x^*) \leq \frac{\varepsilon}{\min(\alpha, 1)}. \tag{31}$$

Proof. As $\tilde{F}^k(\cdot) \leq F(\cdot)$, the stopping criterion implies that

(32)

$$\begin{aligned} F_D(x^k) = \min_x \Big\{F(x) + \frac{1}{2}\big\|x - x^k\big\|^2_D\Big\} &\geq \min_x \Big\{\tilde{F}^k(x) + \frac{1}{2}\big\|x - x^k\big\|^2_D\Big\} \\ &= \tilde{F}^k(z^k_{j_k}) + \frac{1}{2}\big\|z^k_{j_k} - x^k\big\|^2_D \geq F(x^k) - \varepsilon. \end{aligned}$$

Consider two cases.
Case 1: If $F(x^k) - F(x^*) \le \|x^k - x^*\|_D^2$, then (30) with $x = x^k$ yields

$$F_D(x^k) \le F(x^k) - \frac{\big(F(x^k) - F(x^*)\big)^2}{\|x^k - x^*\|_D^2}.$$

Combining this inequality with (32), we conclude that

$$\frac{\big(F(x^k) - F(x^*)\big)^2}{\|x^k - x^*\|_D^2} \le \varepsilon. \tag{33}$$

Substitution of the denominator by the upper estimate $(F(x^k)-F(x^*))/\alpha$ implies (31).
Case 2: $F(x^k) - F(x^*) > \|x^k - x^*\|_D^2$. Then (30) yields

$$F_D(x^k) \le F(x^k) - 2\big(F(x^k) - F(x^*)\big) + \|x^k - x^*\|_D^2.$$

With a view to (32), we obtain

$$2\big(F(x^k) - F(x^*)\big) - \|x^k - x^*\|_D^2 \le \varepsilon,$$

which implies that $F(x^k) - F(x^*) \le \varepsilon$ in this case. □

Lemma 4.4. *Suppose Assumption 4.1 is satisfied. Then at any iteration k we have*

$$F(x^k) - \eta^k \ge \frac{2\varphi(\alpha)}{1 + 2\varphi(\alpha)}\big(F(x^k) - F(x^*)\big).$$

Proof. By Lemma 3.2,

$$F(x^k) - \eta^k \ge F(x^k) - F_D(x^k).$$

To derive a lower bound for the right hand side of the last inequality, we use Assumption 4.1 in (30) with $x = x^k$. We obtain

$$F_D(x^k) \le F(x^k) - \|x^k - x^*\|_D^2\, \varphi(\alpha). \tag{34}$$

By the definition of the Moreau–Yosida regularization, for any optimal solution x^* we have

$$F(x^*) + \frac{1}{2}\|x^* - x^k\|_D^2 \ge F_D(x^k),$$

and thus

$$\|x^k - x^*\|_D^2 \ge 2\big(F_D(x^k) - F(x^*)\big).$$

Substitution to (34) yields

$$F(x^k) - F_D(x^k) \ge 2\big(F_D(x^k) - F(x^*)\big)\varphi(\alpha),$$

which can be manipulated to

$$F(x^k) - F_D(x^k) \ge \frac{2\varphi(\alpha)}{1 + 2\varphi(\alpha)}\big(F(x^k) - F(x^*)\big).$$

This can be combined with the first inequality in the proof, to obtain the desired result. □

In order to estimate the number of iterations of the method needed to achieve the prescribed accuracy, we need to consider two aspects. First, we prove linear rate of convergence between descent steps. Then, we estimate the numbers of null steps between consecutive descent steps.

By employing the estimate of Lemma 4.3, we can address the first aspect. To simplify notation, with no loss of generality, we assume that $\alpha \in (0, 1]$ (otherwise, we would have to replace α with $\bar{\alpha} = \min(\alpha, 1)$ in the considerations below).

Lemma 4.5. *Suppose x^* is the unique minimum point of $F(\cdot)$ and Assumption 4.1 is satisfied. Then at every descent step k, when the update step rule* (8) *is satisfied, we have the inequality:*

$$F(z_{j_k}^k) - F(x^*) \leq (1 - \alpha\beta)\big(F(x^k) - F(x^*)\big). \tag{35}$$

Proof. It follows from the update rule (8) that

$$F(z_{j_k}^k) \leq F(x^k) - \beta\big(F(x^k) - \tilde{F}^k(z_{j_k}^k)\big).$$

Using Lemma 4.3 with $\varepsilon = F(x^k) - \tilde{F}^k(z_{j_k}^k)$, we obtain

$$F(x^k) - F(x^*) \leq \frac{1}{\alpha}\big(F(x^k) - \tilde{F}^k(z_{j_k}^k)\big).$$

Combining these inequalities and simplifying, we conclude that

$$\begin{aligned} F(z_{j_k}^k) &\leq (1 - \beta)F(x^k) + \beta\big(\alpha F(x^*) - \alpha F(x^k) + F(x^k)\big) \\ &= F(x^k) - \alpha\beta\big(F(x^k) - F(x^*)\big). \end{aligned}$$

Subtracting $F(x^*)$ from both sides, we obtain the linear rate (35). □

We now pass to the second issue: the estimation of the number of null steps between two consecutive descent steps. We shall base it on the analysis of the gap $F(x^k) - \eta^k$.

By virtue of Corollary 3.6, all points $\{z_{j_k}^k\}$ generated by the algorithm are uniformly bounded. Since subgradients of finite-valued convex functions are locally bounded, the subgradients of all f_{j_k} are bounded, and thus a constant M exists, such that

$$\big\|s_{j_{k+1}}^k - g_{j_{k+1}}^k\big\|_{D^{-1}}^2 \leq M$$

at all null steps. With no loss of generality, we assume that $\varepsilon \leq (N - 1)M$.

Lemma 4.6. *If a null step is made at iteration k, then*

$$F(x^k) - \eta^{k+1} \leq \gamma\big(F(x^k) - \eta^k\big), \tag{36}$$

where

$$\gamma = 1 - \frac{1}{2}\left(\frac{1 - \beta}{N - 1}\right)^2 \frac{\varepsilon}{M}. \tag{37}$$

Proof. By Lemma 3.1, we have

$$F(x^k) - \eta^{k+1} \le F(x^k) - \eta^k - \frac{1-\beta}{2(N-1)}\bar{\mu}_k v_k. \tag{38}$$

On the other hand,

$$v_k = F(x^k) - \tilde{F}^k(z^k_{j_k}) = F(x^k) - \eta^k + \frac{1}{2}\big\|z^k_{j_k} - x^k\big\|^2_D \ge F(x^k) - \eta^k. \tag{39}$$

Combining the last two inequalities, we conclude that

$$\begin{aligned} F(x^k) - \eta^{k+1} &\le F(x^k) - \eta^k - \frac{1-\beta}{2(N-1)}\bar{\mu}_k\big(F(x^k) - \eta^k\big) \\ &= \left(1 - \frac{1-\beta}{2(N-1)}\bar{\mu}_k\right)\big(F(x^k) - \eta^k\big). \end{aligned} \tag{40}$$

Consider the definition (12) of $\bar{\mu}_k$ in Lemma 3.1. If $\bar{\mu}_k = 1$, then $1 - \frac{1-\beta}{2(N-1)}\bar{\mu}_k$ is no greater than the bound (37), because $\varepsilon \le (N-1)M$. Otherwise, $\bar{\mu}_k$ is given by the second case in (12). Since the algorithm does not stop, we have $v_k > \varepsilon$, and thus

$$\bar{\mu}_k \ge \frac{(1-\beta)\varepsilon}{(N-1)M}.$$

Substitution to (40) yields (37). □

Let $x^{(\ell-1)}, x^{(\ell)}, x^{(\ell+1)}$ be three consecutive proximal centers in the algorithm ($\ell \le 2$). We want to bound the number of iterations with the proximal center $x^{(\ell)}$. To this end, we bound two quantities:

1. The optimal objective value of the *first* subproblem with proximal center $x^{(\ell)}$, whose iteration number we denote by $k(\ell)$:

$$\eta^{k(\ell)} = \min f_{j_{k(\ell)}}(x) + \sum_{i \ne j_{k(\ell)}} \tilde{f}_i^{k(\ell)}(x) + \frac{1}{2}\big\|x - x^{(\ell)}\big\|^2_D. \tag{41}$$

We need an upper bound for $F(x^{(\ell)}) - \eta^{k(\ell)}$.

2. The optimal objective value of the *last* subproblem with proximal center $x^{(\ell)}$, occurring at iteration $k'(\ell) = k(\ell+1) - 1$:

$$\eta^{k'(\ell)} = \min f_{j_{k'(\ell)}}(x) + \sum_{i \ne j_{k'(\ell)}} \tilde{f}_i^{k'(\ell)}(x) + \frac{1}{2}\big\|x - x^{(\ell)}\big\|^2_D. \tag{42}$$

We need an upper bound for $F(x^{(\ell)}) - \eta^{k'(\ell)}$ which implies the update rule (8).

In the following we discuss each issue separately.

Recall that according to the algorithm, $x^{(\ell)}$ is the optimal solution of the last subproblem with proximal center $x^{(\ell-1)}$. Let $f_{j_{k(\ell)-1}}$ be the non-linearized component function of the last subproblem with proximal center $x^{(\ell-1)}$, whose optimal solution is $x^{(\ell)}$. The optimal value of the subproblem (5) is

$$\eta^{k(\ell)-1} = f_{j_{k(\ell)-1}}(x^\ell) + \sum_{i \ne j_{k(\ell)-1}} \tilde{f}_i^{k(\ell)-1}(x^{(\ell)}) + \frac{1}{2}\big\|x^{(\ell)} - x^{(\ell-1)}\big\|^2_D. \tag{43}$$

Lemma 4.7. *If a descent step is made at iteration $k(\ell)-1$, then*

$$F(x^{(\ell)}) - \eta^{k(\ell)} \leq \frac{3}{2\beta}\big(F(x^{(\ell-1)}) - F(x^{(\ell)})\big). \tag{44}$$

Proof. The left inequality in (28) yields

$$\eta^{k(\ell)} \geq \eta^{k(\ell)-1} - \big\|x^{(\ell)} - x^{(\ell-1)}\big\|_D^2.$$

Since $F(x^{(\ell)}) \leq F(x^{(\ell-1)})$, we obtain

$$F(x^{(\ell)}) - \eta^{k(\ell)} \leq F(x^{(\ell-1)}) - \eta^{k(\ell)-1} + \big\|x^{(\ell)} - x^{(\ell-1)}\big\|_D^2.$$

As iteration $k(\ell)-1$ is a descent step, the update rule (8) holds. Thus

$$\begin{aligned} F(x^{(\ell-1)}) - \eta^{k(\ell)-1} &= \Big[F(x^{(\ell-1)}) - \tilde{F}^{k(\ell)-1}(x^{(\ell)})\Big] - \frac{1}{2}\big\|x^{(\ell)} - x^{(\ell-1)}\big\|_D^2 \\ &\leq \frac{1}{\beta}\Big[F(x^{(\ell-1)}) - F(x^{(\ell)})\Big] - \frac{1}{2}\big\|x^{(\ell)} - x^{(\ell-1)}\big\|_D^2. \end{aligned}$$

Combining the last two inequalities we obtain

$$F(x^{(\ell)}) - \eta^{k(\ell)} \leq \frac{1}{\beta}\big(F(x^{(\ell-1)}) - F(x^{(\ell)})\big) + \frac{1}{2}\big\|x^{(\ell)} - x^{(\ell-1)}\big\|_D^2.$$

The right inequality in (28) can be now used to substitute $\big\|x^{(\ell)} - x^{(\ell-1)}\big\|_D^2$ on the right hand side to obtain (44). □

We can now integrate our results.

Applying Lemma 4.4, we obtain the following inequality at *every* null step with prox center $x^{(\ell)}$:

$$\begin{aligned} F(x^{(\ell)}) - \eta^{k} &\geq \frac{2\varphi(\alpha)}{1+2\varphi(\alpha)}\big(F(x^{(\ell)}) - F(x^*)\big) \\ &\geq \frac{2\varphi(\alpha)}{1+2\varphi(\alpha)}\big(F(x^{(\ell)}) - F(x^{(\ell+1)})\big). \end{aligned} \tag{45}$$

From Lemma 4.7 we know that for $\ell \geq 2$ the initial value of the left hand side (immediately after the previous descent step) is bounded from above by the following expression:

$$F(x^{(\ell)}) - \eta^{k(\ell)} \leq \frac{3}{2\beta}\big(F(x^{(\ell-1)}) - F(x^{(\ell)})\big). \tag{46}$$

Lemma 4.6 established a linear rate of decrease of the left hand side. Therefore, the number n_ℓ of null steps with proximal center $x^{(\ell)}$, if it is positive, satisfies the inequality:

$$\frac{3}{2\beta}\big(F(x^{(\ell-1)}) - F(x^{(\ell)})\big)\gamma^{n_\ell - 1} \geq \frac{2\varphi(\alpha)}{1+2\varphi(\alpha)}\big(F(x^{(\ell)}) - F(x^{(\ell+1)})\big).$$

Consequently, for $\ell \geq 2$ we obtain the following upper bound on the number of null steps:

$$n_\ell \leq 1 + \frac{1}{\ln(\gamma)} \ln\left(\frac{4\beta\varphi(\alpha)}{3(1+2\varphi(\alpha))} \frac{F(x^{(\ell)}) - F(x^{(\ell+1)})}{F(x^{(\ell-1)}) - F(x^{(\ell)})}\right). \tag{47}$$

If the number n_ℓ of null steps is zero, inequality (35) yields

$$\frac{F(x^{(\ell)}) - F(x^{(\ell+1)})}{F(x^{(\ell-1)}) - F(x^{(\ell)})} \leq \frac{F(x^{(\ell)}) - F(x^*)}{F(x^{(\ell-1)}) - F(x^*) - \big(F(x^{(\ell)}) - F(x^*)\big)} \leq \frac{1}{\frac{1}{1-\alpha\beta} - 1}.$$

Elementary calculations prove that both logarithms on the right hand side of (47) are negative, and thus inequality (47) is satisfied in this case as well.

Suppose there are L proximal centers appearing throughout the algorithm: $x^{(1)}$, $x^{(2)}$, …, $x^{(L)}$. They divide the progress of the algorithm into L series of null steps. For the first series, similarly to the analysis above, we use (45) and Lemma 4.6 to obtain the estimate

$$n_1 \leq 1 + \frac{1}{\ln(\gamma)} \ln\left(\frac{2\varphi(\alpha)}{1+2\varphi(\alpha)} \frac{F(x^{(1)}) - F(x^{(2)})}{F(x^{(1)}) - \eta^1}\right). \tag{48}$$

For the last series, we observe that the inequality $F(x^{(\ell)}) - \eta^k \geq \varepsilon/2$ has to hold at each null step at which the stopping test was not satisfied. We use it instead of (45) and we obtain

$$n_L \leq 1 + \frac{1}{\ln(\gamma)} \ln\left(\frac{\beta}{3} \frac{\varepsilon}{F(x^{(L-1)}) - F(x^{(L)})}\right). \tag{49}$$

We aggregate the total number of null steps for different proximal centers throughout the algorithm and we obtain the following bound:

$$\sum_{\ell=1}^{L} n_\ell = \frac{L-1}{\ln(\gamma)} \left[\ln\left(\frac{2\varphi(\alpha)}{1+2\varphi(\alpha)}\right) + \ln\left(\frac{\beta}{3}\right) + \frac{1}{L-1} \ln\left(\frac{\varepsilon}{F(x^{(1)}) - \eta^1}\right)\right] + L \tag{50}$$

Let us recall the definition of γ in (37), and denote

$$C = \frac{1}{2}\left(\frac{1-\beta}{N-1}\right)^2 \frac{1}{M},$$

so that $\gamma = 1 - \varepsilon C$. Since $\ln(1-\varepsilon C) < -\varepsilon C$, we derive the following inequality for the number of null steps:

$$\sum_{\ell=1}^{L} n_\ell \leq \frac{L-1}{-\varepsilon C} \left[\ln\left(\frac{2\varphi(\alpha)}{1+2\varphi(\alpha)}\right) + \ln\left(\frac{\beta}{3}\right) + \frac{1}{L-1} \ln\left(\frac{\varepsilon}{F(x^{(1)}) - \eta^1}\right)\right] + L. \tag{51}$$

Let us now derive an upper bound on the number L of descent steps. By virtue of (7) and (8), descent steps are made only if

$$F(x^k) - F(x^*) \geq \beta\varepsilon;$$

otherwise, the method must stop. To explain it more specifically, if $F(x^k) - F(x^*) \leq \beta\varepsilon$, then $F(x^k) - F(z_{j_k}^k) \leq \beta\varepsilon$. If a descent step were made, $F(z_{j_k}^k) \leq F(x^k) - \beta v_k$. Then $\beta v_k \leq \beta\varepsilon$. Since $v_k \leq \varepsilon$, the algorithm would have already stopped, which contradicts our assumption. It follows from Lemma 4.5, that

$$(1-\alpha\beta)^{L-1}\big(F(x^{(1)}) - F(x^*)\big) \geq \beta\varepsilon.$$

Therefore,

$$L \leq 1 + \frac{\ln(\beta\varepsilon) - \ln\big(F(x^1) - F(x^*)\big)}{\ln(1-\alpha\beta)}. \tag{52}$$

As a result, we have the final bound for the total number of descent and null steps:

(53)

$$\begin{aligned} L + \sum_{\ell=1}^{L} n_\ell \leq \frac{1}{\varepsilon C \ln(1-\alpha\beta)} \ln \frac{F(x^{(1)}) - F(x^*)}{\beta\varepsilon} \left[\ln\left(\frac{2\varphi(\alpha)}{1+2\varphi(\alpha)}\right) + \ln\left(\frac{\beta}{3}\right) \right] \\ + \frac{1}{\varepsilon C} \ln\left(\frac{F(x^{(1)}) - \eta^1}{\varepsilon}\right) + 2\frac{\ln(\beta\varepsilon) - \ln\big(F(x^{(1)}) - F(x^*)\big)}{\ln(1-\alpha\beta)} + 2. \end{aligned}$$

Therefore, in order to achieve precision ε, the number of steps needed is of order

$$L + \sum_{\ell=1}^{L} n_\ell \sim \mathcal{O}\left(\frac{1}{\varepsilon}\ln\left(\frac{1}{\varepsilon}\right)\right).$$

This is almost equivalent to saying that given the number of iterations k, the precision of the solution is approximately $\mathcal{O}(1/k)$.

5 Application to structured regularized regression problems

In many areas in data mining and machine learning, such as computer vision and compressed sensing, the resulting optimization models consist of a convex loss function and multiple convex regularization functions, called the *composite prior models* in [HZLM11]. For example, in compressed sensing, the linear combination of the total variation (TV) norm and L_1 norm is a popular regularizer in recovering Magnetic Resonance (MR) images. Formally, the models are formulated as follows:

$$\min_{x\in\mathbb{R}^n} F(x) = f(x) + \sum_{i=1}^{N} h_i(B_i x), \tag{54}$$

where f is the loss function to measure the goodness-of-fit of the data, while the functions h_i are regularization terms. All the functions are convex but not necessarily smooth.

The SLIN algorithm introduced in our paper can be directly applied to solve the general problem (54). It can be further specialized to take advantage of additional features of the functions involved. In the following subsection we discuss one such specialization.

5.1 Fused lasso regularization problem

The problem is defined as follows:

$$\min_x \frac{1}{2}\|b - Ax\|_2^2 + \lambda_1 \|x\|_1 + \lambda_2 \sum_{j=1}^{p-1} |x_{j+1} - x_j|, \tag{55}$$

where A is an $m \times n$ matrix, and $\lambda_1, \lambda_2 > 0$ are fixed parameters. This model contains two regularization terms: the *lasso* penalty $h_1(x) = \lambda_1 \|x\|_1$, and the *fused lasso* penalty $h_2(x) = \lambda_2 \sum_j |x_{j+1} - x_j|$. We name the first function as $f(x) := \frac{1}{2}\|b - Ax\|_2^2$. In models with a quadratic loss function, we found it convenient to use the matrix $D = \text{diag}(A^T A)$ in the proximal term of the method [LPR14, Sec. 3].

In order to solve each subproblem, we need the gradient of $f(\cdot)$ and subgradients of the regularization functions, which are readily available. Our method requires their explicit calculation at the initial iteration only; at later iterations they are obtained implicitly, as described in Step 1 of the algorithm.

With these, we can solve each subproblem iteratively.

The f-subproblem. Skipping the constants, the f-subproblem has the form:

$$\min_x \frac{1}{2}\|b - Ax\|_2^2 + g_{h_1}^T x + g_{h_2}^T x + \frac{1}{2}\|x - x^k\|_D^2. \tag{56}$$

This is a unconstrained quadratic optimization problem and its optimal solution can be obtained by solving the following linear system of equations:

$$(A^T A + D)x = A^T b - g_{h_1} - g_{h_2} + D^T x^k.$$

It can be very efficiently solved by the preconditioned conjugate gradient method with preconditioner D, as discussed in [LPR14, Sec. 3], because the condition index of the system is uniformly bounded. Only matrix–vector multiplications are involved, facilitating the use of a sparse structure of A. After the solution is obtained, the gradient of $f(x)$ and its linearization can be determined by Step 1 of the SLIN algorithm.

The h_1-subproblem. The subproblem is defined as follows (ignoring the constants):

$$\min_x g_f^T x + \lambda_1 \|x\|_1 + g_{h_2}^T x + \frac{1}{2}\|x - x^k\|_D^2. \tag{57}$$

This problem is separable in the decision variables, with the following closed-form solution:

$$(x_{h_1})_i = \text{sgn}(\tau_i) \max\left(0, |\tau_i| - \frac{\lambda_1}{d_i}\right), \quad i = 1, \ldots, n.$$

Here $\tau_i = x_i^k - \frac{(g_f)_i + (g_{h_2})_i}{d_i}$.

The solution of the h_1-subproblem gives a new subgradient of h_1 at the minimal point.

The h_2-subproblem. The subproblem is defined as follows (ignoring the constants):

$$\min_x g_f^T x + g_{h_1}^T x + \lambda_2 \sum_{j=1}^{p-1} |x_{j+1} - x_j| + \frac{1}{2}\|x - x^k\|_D^2.$$

Exactly as described in [LPR14], this problem can be equivalently formulated as a constrained optimization problem:

$$\min_{x,z} g_f^T x + g_{h_1}^T x + \lambda_2 \|z\|_1 + \frac{1}{2}\|x - x^k\|_D^2, \text{ subject to } Rx = z,$$

with an $(n-1)\times n$ matrix R representing the system $z_j = x_{j+1} - x_j$, $j = 1, \ldots, n-1$. The Lagrangian of problem (19) has the form

$$L(x, z, \mu) = g_f^T x + g_{h_1}^T x + \lambda_2 \|z\|_1 + \mu^T (Rx - z) + \frac{1}{2}\|x - x^k\|_D^2,$$

where μ is the dual variable. The minimum of the Lagrangian with respect to z is finite if and only if $\|\mu\|_\infty \leq \lambda_2$. Under this condition, the minimum value of the z-terms is zero and we can eliminate them from the Lagrangian. We arrive to its reduced form,

$$\hat{L}(x, \mu) = g_f^T x + g_{h_1}^T x + \mu^T R x + \frac{1}{2}\|x - x^k\|_D^2.$$

To calculate the dual function, we minimize $\hat{L}(x, \mu)$ with respect to x. After elementary calculations, we obtain the solution

$$\tilde{x}_{h_2} = x^k - D^{-1}(g_f + g_{h_1} + R^T \mu).$$

Substituting it back to the Lagrangian, we obtain the following dual problem:

$$\max_{\mu} -\frac{1}{2}\mu^T R D^{-1} R^T \mu + \mu^T R(x^k - D^{-1} g_f - D^{-1} g_h), \text{ subject to } \|\mu\|_\infty \leq \lambda_2.$$

This problem can be treated as a box-constrained quadratic programming problem, for which many efficient algorithms are available, for example coordinate-wise optimization [LPR14, Sec. 4]. Due to the structure of R, the computational effort per iteration is linear in the problem dimension.

6 Overlapping group lasso problem

We consider the following problem

$$\min_x \frac{1}{2K\lambda}\|b - Ax\|_2^2 + \sum_{j=1}^{K} d_j \|x_{\mathcal{G}_j}\|_2 \tag{58}$$

where $A \in R^{m\times n}$. This model contains the first function as $f(x) := \frac{1}{2K\lambda}||b - Ax||_2^2$ where parameter $\lambda > 0$ and number of groups K are pre-specified parameters. The second part is a sum of regularization terms, each penalty function as $h_j(x) = d_j||x_{\mathcal{G}_j}||_2$ where the weights $d_j > 0$ are known parameters. $\mathcal{G}_j \subseteq \{1, \ldots, p\}$ is the index set of a group of variables and $x_{\mathcal{G}_j}$ denotes the subvector of x with coordinates in $\mathcal{G}_j$. This group regularizer has been proven useful in high-dimensional statistics with the capability of selecting meaningful groups of features. The groups could overlap as needed.

As the quadratic term has a coefficient of $\frac{1}{2K\lambda}$, the diagonal matrix D in the proximal term of the method is set to $D = \frac{1}{K\lambda}\text{diag}(A^TA)$.

The f-subproblem. The f-subproblem has the form:

$$\min_x \frac{1}{2K\lambda}\|b - Ax\|_2^2 + \sum_{j=1}^{K} g_j^T x + \frac{1}{2}\|x - x^k\|_D^2.$$

It has the same structure as the f-subproblem of the general structured fused lasso example, and can be solved in the same way; just the matrix D is different.

The h_j-subproblem. The h_j-subproblem is defined as follows (ignoring the constants):

$$\min d_j\|x_{\mathcal{G}_j}\|_2 + \langle s, x\rangle + \frac{1}{2}\|x - x^k\|_D^2. \tag{59}$$

where $s = g_f + \sum_{j' \neq j} g_{h_{j'}}$; with g_f denoting a subgradient of the function f, and $g_{h_{j'}}$ the subgradients of $h_{j'}$ used in (5). To simplify notation, from now on we write $\mathcal{G}$ for $\mathcal{G}_j$.

The decision variables that are outside of the current group $\mathcal{G}$, which we denote $x_{-\mathcal{G}}$, have the following closed-form solution:

$$x_{-\mathcal{G}} = x^k_{-\mathcal{G}} - D^{-1}_{-\mathcal{G}} s_{-\mathcal{G}}.$$

The variables in the current group $\mathcal{G}$ can be calculated as follows. If $x_{\mathcal{G}} \neq 0$, the necessary and sufficient optimality condition for (59) is the following equation:

$$\frac{d_j x_{\mathcal{G}}}{\|x_{\mathcal{G}}\|_2} + s_{\mathcal{G}} + D_{\mathcal{G}}(x_{\mathcal{G}} - x^k_{\mathcal{G}}) = 0. \tag{60}$$

We denote

$$\frac{d_j}{\|x_{\mathcal{G}}\|_2} = \kappa, \tag{61}$$

This leads to

$$x_i = \frac{D_{ii}x_i^k - (s)_i}{\kappa + D_{ii}}, \quad i \in \mathcal{G}. \tag{62}$$

Substituting into (61), after simple manipulations, we obtain the following equation for κ:

$$\sum_{i \in \mathcal{G}_j} \left(\frac{D_{ii}x_i^k - (s)_i}{1 + \frac{D_{ii}}{\kappa}}\right)^2 = d_j^2. \tag{63}$$

Since the left hand side of this equation is an increasing function of κ, we can easily solve it by bisection, if a solution exists. If the columns if A are normalized, then all $D_{ii} = 1$, and equation (63) can be solved in closed form.

Letting $\kappa \to \infty$ on the left hand side, we obtain the condition for the existence of a solution of (63):

$$\sum_{i \in \mathcal{G}_j} \left(D_{ii}x_i^k - (s)_i\right)^2 > d_j^2. \tag{64}$$

If inequality (64) is satisfied, κ can be found by bisection and $x_{\mathcal{G}}$ follows from (62). If (64) is not satisfied, the only possibility is that the optimal solution of (59) is $x_{\mathcal{G}} = 0$.

7 Numerical Results

In this section, we present some experimental results for problems (55) and (58). All these studies are performed on an 1.8 GHZ, 4GB RAM computer using MATLAB.

7.1 Fused lasso experiments

In Tables 1 and 2, we evaluate SLIN against three its modifications, in order to assess the usefulness of the main features of the method: the dynamic selection of the block to be optimized, and the sufficient improvement test. The first modification processes the blocks in a fixed order and performs the improvement test after every block, to decide whether to change the current value of x^k. In a way, it is a direct extension of the alternating linearization method (ALIN) of [LPR14], and is labeled as such in the tables. The second modification processes the blocks in a fixed order and updates x^k after each cycle of 3 blocks. In the case of two blocks, it would correspond to the Douglas–Rachford operator splitting method of [LM79]; see [LPR14]. We use the name Douglas–Rachford in the tables. The third modification processes the blocks in a fixed order and carries out the update of x^k after each block. In the case of two blocks, it would correspond to the Peaceman–Rachford operator splitting method of [LM79], as explained in [LPR14], and we use this name in the tables.

We report the average performance of all 4 versions in cases when $m > n$ $(m = 1000, n = 300)$ and when $m < n$ $(m = 300, n = 1000)$, with different tolerance and regularization parameters. We run the experiments 10 times with different samples of the matrix A and report the average results and their standard deviation. In the tables, "Iterations" denotes the average iteration number with 1000 as the default maximum iteration number; "Time" denotes the average CPU run time in seconds; "Relative Error" is defined as the relative difference between the optimal value (obtained by MATLAB "fminunc" function) and those obtained by the SLIN, ALIN, and operator splitting methods, respectively. "MAXITER" and "NA" indicate that the algorithm did not converge in the pre-specified number of 1000 iterations. We observe that to obtain the same accuracy, SLIN requires significantly fewer iterations and less CPU time than the other versions. The versions without sufficient progress test (naïve extensions of operator splitting methods) did not converge in many cases. The version with the fixed block order (naïve extension of the alternating linearization method), although convergent, was always significantly slower.

In Table 3, we report performance of SLIN with different values of the parameter β. Based on that, we use $\beta = 0.5$ in all our further experiments.

7.2 Overlapping group lasso experiments

We compared our method with existing algorithms in the tree-structured, fixed order, and random order cases. It shows that SLIN does not care how the groups are partitioned and is applicable to arbitrary group-splitting cases.

Parameters	Method	Iterations	Time	Relative Error
$tol = 10^{-3}$	Selective Linearization	55(28.75)	0.99(0.47)	0.0156(0.0043)
$\lambda = 1\tau$	Alternating Linearization	66(31.04)	1.10 (0.99)	0.0161(0.0048)
	Douglas-Rachford	323(78.57)	5.53 (0.96)	0.0161(0.0046)
	Peaceman-Rachford	MAXITER(0.00)	16.93 (0.23)	0.9191(0.2123)
$tol = 10^{-7}$	Selective Linearization	46 (25.79)	1.00 (0.41)	0.0154(0.0075)
$\lambda = 1\tau$	Alternating Linearization	173 (31.30)	2.98 (0.60)	0.0166(0.0078)
	Douglas-Rachford	246(42.57)	4.32 (0.73)	0.0155(0.0076)
	Peaceman-Rachford	MAXITER (0.00)	16.75 (0.22)	NA(NA)
$tol = 10^{-3}$	Selective Linearization	79(2.30)	1.60 (0.07)	0.0004(0.0000)
$\lambda = 10^{-2}\tau$	Alternating Linearization	91 (3.05)	1.64 (0.06)	0.0004(0.0000)
	Douglas-Rachford	205(10.17)	3.45 (0.21)	0.0004(0.0000)
	Peaceman-Rachford	MAXITER (0.00)	16.57 (0.11)	NA(NA)
$tol = 10^{-7}$	Selective Linearization	177(20.68)	3.06(0.88)	0.0028(0.0019)
$\lambda = 10^{-2}\tau$	Alternating Linearization	450(50.74)	8.01(4.06)	0.0028(0.0019)
	Douglas-Rachford	573(59.81)	9.91 (0.52)	0.0028(0.0019)
	Peaceman-Rachford	MAXITER (0.00)	16.60(0.15)	NA(NA)

Table 1: Comparison of methods on problems with $m = 1000$ and $n = 300$.

7.2.1 Tree-structured overlapping groups

In a tree-structured overlapping group lasso problem, described in [JMOB11], the groups correspond to nodes of a tree. Thus, for any two groups, either one is a subset of the other, or they are disjoint. The design matrix and input vector are centered and normalized to have unit ℓ_2-norms. We conduct the speed comparisons between our approach and FISTA [JMOB11]. From Table 4 we can see that the SLIN algorithm is faster in terms of both iteration number and computational time.

7.2.2 Fixed order overlapping groups

We simulate data for a univariate linear regression model with an overlapping group structure. The entries are sampled from i.i.d. normal distributions, $x_j = (-1)^j \exp(-(j-1)/100)$, and $b = Ax + \varepsilon$, with the noise ε sampled from the standard normal distribution. Assuming that the inputs are ordered, we define a sequence of K groups of 100 adjacent inputs with an overlap of 10 variables between two successive groups, so that

$$G = \{\{1, \ldots, 100\}, \{91, \ldots, 190\}, \ldots, \{n - 99, \ldots, p\}\}, \tag{65}$$

where $n = 90K + 10$. We adopt uniform weights $d_j = 1/K$ and set $\lambda = K/5$.

To demonstrate the efficiency and scalability of the SLIN algorithm, we compared SLIN with several specialized methods for overlapping group lasso problems: PDMM of [CDZ15, WBL14], sADMM or Jacobian ADMM of [DLPY13], PA-APG of [Yu13] and S-APG of [CLK$^+$12]. All experiments were run sequentially, that is, no parallel

Parameters	Method	Iterations	Time	Relative Error
$tol = 10^{-3}$	Selective Linearization	83(14.15)	11.26(3.01)	0.0393(0.0103)
$\lambda = 1\tau$	Alternating Linearization	425(13.96)	55.97(4.87)	0.0456(0.0126)
	Douglas-Rachford	MAXITER(0.00)	130.39 (6.42)	0.0355(0.0096)
	Peaceman-Rachford	MAXITER(0.00)	131.19 (8.48)	NA(NA)
$tol = 10^{-7}$	Selective Linearization	257 (7.97)	36.76(13.93)	0.0165(0.0059)
$\lambda = 1\tau$	Alternating Linearization	419(8.23)	59.86(1.99)	0.0240(0.0083)
	Douglas-Rachford	MAXITER(0.00)	144.31 (18.14)	0.0187(0.0059)
	Peaceman-Rachford	MAXITER(0.00)	143.93 (11.56)	0.0118(0.0063)
$tol = 10^{-3}$	Selective Linearization	210(27.96)	27.30(18.07)	0.0483(0.0160)
$\lambda = 10^{-2}\tau$	Alternating Linearization	290 (34.73)	39.86(28.05)	0.0484(0.0160)
	Douglas-Rachford	291(53.85)	39.51 (32.29)	0.0484(0.0160)
	Peaceman-Rachford	MAXITER(0.00)	139.38(18.88)	NA(NA)
$tol = 10^{-7}$	Selective Linearization	404 (45.45)	49.51(8.04)	0.0550(0.0181)
$\lambda = 10^{-2}\tau$	Alternating Linearization	780(63.45)	107.23(5.01)	0.0550(0.0181)
	Douglas-Rachford	891(53.07)	106.94 (12.29)	0.0550(0.0181)
	Peaceman-Rachford	MAXITER(0.00)	133.34(7.81)	NA(NA)

Table 2: Comparison of methods on problems with $m = 300$ and $n = 1000$

β	Iterations	Time	Relative Error
0.2	76(25.61)	0.82(0.95)	0.0200(0.0100)
0.5	28 (24.70)	0.39 (0.64)	0.0080(0.0037)
0.8	74(33.72)	1.08 (0.63)	0.0400(0.0200)

Table 3: The effect of different values of β in the SLIN algorithm for the problem with $n = 1000$ and $p = 300$.

Parameters	Methods	Iter	Time
$m = 100, n = 10$	SLIN	11.60 (0.70)	0.0897(0.0035)
$K = 8$	FISTA	25.20(2.85)	0.1385(0.0138)

Table 4: Comparison of SLIN and FISTA on tree-structured overlapping group lasso problem.

processing features were exploited. We run the experiments 10 times with different samples of the matrix A; we report the average results.

Figure 1 plots the convergence of the objective function values *versus* the number of iterations, for the number of groups $K = 100$. For the dual methods PDMM and

sADMM, we report the values of the augmented Lagrangian. They go from super optimal (because the iterates are infeasible) and converge to the optimal value. The SLIN algorithm is the fastest, in terms of iterations while PDMM is the second, and sADMM the third. The two accelerated methods, PA-APG and S-APG, are the slowest in these tests.

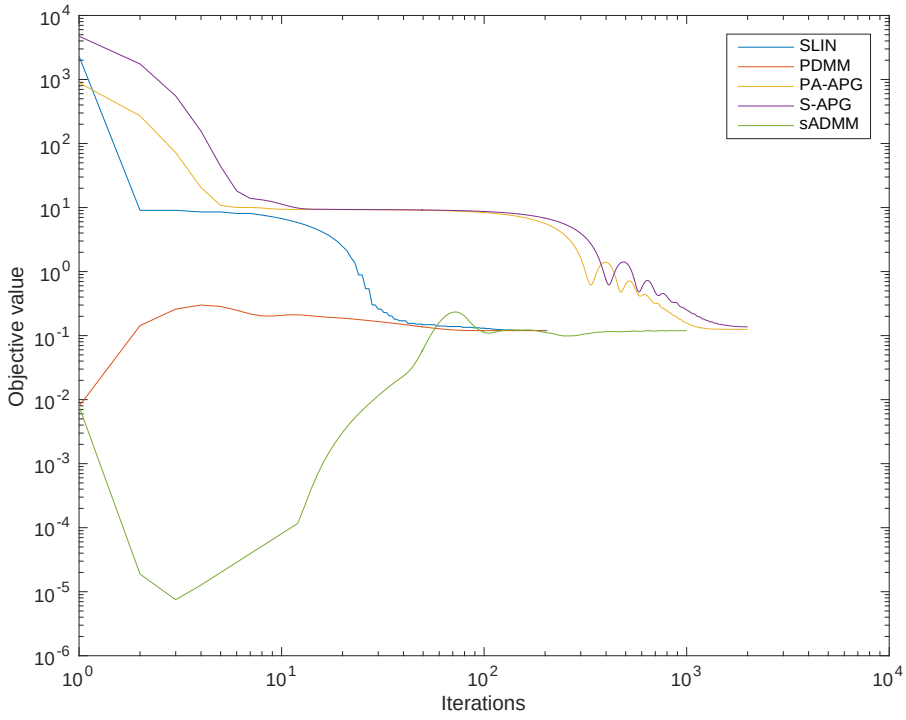

Figure 1: Comparison of SLIN and other algorithms on an overlapping group lasso problem.

7.2.3 Randomly overlapping groups

In the next stage, we conducted additional comparisons between SLIN and PDMM on group lasso problems with randomized overlapping, which do not exhibit the regular group structure specified in (65).

This type of problem arises in applications such as bioinformatics [3], where one uses prior information to model potential overlapping of groups of variables. For example, in high throughput gene expression analysis, the number of parameters to be estimated is much greater than the sample size. One often utilizes information including gene ontology to define group overlaps among genes, thereby achieving structured regularization [VRMV14]. The resulting overlaps are "arbitrary" (depending on the specific

gene ontology) and more complex than the systematic overlapping example described in (65). We generated test cases in which the indices in each of the 100 groups were assigned to the n locations. As a result, the number of overlapping variables between the groups was random, and multiple group membership was possible. The performance of the two methods on randomized overlapping group lasso problems is summarized in Tables 4 and 5.

For fair comparison of the methods, we run PDMM on each instance of the problem. PDMM was set to run to "tol" $= 10^{-4}$ or 2,000 iterations, whichever came first. We set the tuning parameters $d_g = 0.01/K$, and $0.02/K$, respectively. Then SLIN was set to run until the objective function values obtained were as good as that of PDMM. We run the experiments 10 times with different samples of the randomly generated groups; in Tables 5 and 6 we report the average results and their standard deviations. In all cases, the number of iterations of SLIN is much smaller than that of PDMM. In the determined cases, where $m = 1000$ and $n = 800$, the running time of SLIN is usually better than that of PDMM. In the under-determined cases, where $m = 500$ and $n = 600$, the running time of SLIN is slightly worse than that of PDMM.

In summary, we can conclude that SLIN is a highly efficient and reliable general-purpose method for multi-block optimization of convex nonsmooth functions. It successfully competes with dedicated methods for special classes of problems.

Parameters	Methods	Iter	Time
$m = 1000, n = 800$	SLIN	280 (15.82)	5.42(0.36)
$K = 80, d_g = 0.01/K$	PDMM	638(21.14)	5.48(0.10)
$m = 1000, n = 800$	SLIN	308 (16.11)	5.38(0.43)
$K = 90, d_g = 0.01/K$	PDMM	836(29.33)	7.19(0.39)
$m = 1000, n = 800$	SLIN	331 (13.05)	5.89(0.28)
$K = 100, d_g = 0.01/K$	PDMM	991 (75.51)	8.91(0.81)
$m = 1000, n = 800$	SLIN	306 (11.79)	5.53(0.23)
$K = 80, d_g = 0.02/K$	PDMM	560(59.47)	4.75(0.45)
$m = 1000, n = 800$	SLIN	330 (11.19)	5.63(0.22)
$K = 90, d_g = 0.02/K$	PDMM	620(38.33)	5.65(0.91)
$m = 1000, n = 800$	SLIN	364 (13.77)	6.09(0.31)
$K = 100, d_g = 0.02/K$	PDMM	741(58.06)	4.14 (0.47)

Table 5: Comparison SLIN and PDMM in solving an overlapping group lasso of randomly generated groups. Determined cases with $m = 1000$ and $n = 800$.

References

[BC11] H. H. Bauschke and P. L. Combettes. *Convex Analysis and Monotone Operator Theory in Hilbert Spaces*. CMS Books in Mathematics/Ouvrages

Parameters	Methods	Iter	Time
$m = 500, n = 600$	SLIN	1280 (39.73)	17.75(0.60)
$K = 80, d_g = 0.01/K$	PDMM	1453(157.07)	12.01(1.03)
$m = 500, n = 600$	SLIN	1119 (57.76)	16.61(1.00)
$K = 90, d_g = 0.01/K$	PDMM	1566(61.63)	13.88(1.63)
$m = 500, n = 600$	SLIN	973 (39.46)	13.72(1.241)
$K = 100, d_g = 0.01/K$	PDMM	1753 (122.27)	16.75(2.64)
$m = 500, n = 600$	SLIN	792 (36.58)	9.79(0.73)
$K = 80, d_g = 0.02/K$	PDMM	968(39.82)	7.53(0.54)
$m = 500, n = 600$	SLIN	722 (23.46)	8.87(0.65)
$K = 90, d_g = 0.02/K$	PDMM	1170(102.16)	9.57(1.00)
$m = 500, n = 600$	SLIN	683 (15.66)	8.01(0.41)
$K = 100, d_g = 0.02/K$	PDMM	1208(70.30)	10.81 (1.21)

Table 6: Comparison SLIN and PDMM in solving an overlapping group lasso of randomly generated groups. Underdetermined cases with $m = 500$ and $n = 600$

de Mathématiques de la SMC. Springer, New York, 2011.

[Ber15] D. P. Bertsekas. Incremental aggregated proximal and augmented lagrangian algorithms. Technical Report Report LIDS-3176, Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, Mass., 2015.

[BPC+10] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends in Machine Learning*, 3(1):1–122, 2010.

[CDZ15] N. Chatzipanagiotis, D. Dentcheva, and M. M. Zavlanos. An augmented lagrangian method for distributed optimization. *Mathematical Programming*, 152(1-2):405–434, 2015.

[CLK+12] X. Chen, Q. Lin, S. Kim, J. G. Carbonell, and E. P. Xing. Smoothing proximal gradient method for general structured sparse regression. *The Annals of Applied Statistics*, 6(2):719–752, 2012.

[Com09] P. L. Combettes. Iterative construction of the resolvent of a sum of maximal monotone operators. *J. Convex Anal.*, 16(3-4):727–748, 2009.

[Con13] L. Condat. A primal–dual splitting method for convex optimization involving lipschitzian, proximable and linear composite terms. *Journal of Optimization Theory and Applications*, 158(2):460–479, 2013.

[CP11] P. L. Combettes and J.-C. Pesquet. Proximal splitting methods in signal processing. In *Fixed-Point Algorithms for Inverse Problems in Science and Engineering*, Springer Optimization and Its Applications, pages 185–212. Springer, 2011.

[DLPY13] W. Deng, M.-J. Lai, Z. Peng, and W. Yin. Parallel multi-block admm with o (1/k) convergence. Technical report, 2013.

[DR56] J. Douglas and H. H. Rachford. On the numerical solution of heat conduction problems in two and three space variables. *Trans. Amer. Math. Soc.*, 82:421–439, 1956.

[EB92] J. Eckstein and D. P. Bertsekas. On the Douglas-Rachford splitting method and the proximal point algorithm for maximal monotone operators. *Math. Programming*, 55(3, Ser. A):293–318, 1992.

[GLT89] R. Glowinski and P. Le Tallec. *Augmented Lagrangian and operator-splitting methods in nonlinear mechanics*, volume 9. SIAM, 1989.

[GM75] R. Glowinski and A. Marroco. Sur l'approximation, par éléments finis d'ordre un, et la résolution, par pénalisation-dualité d'une classe de problèmes de dirichlet non linéaires. *Revue française d'automatique, informatique, recherche opérationnelle. Analyse numérique*, 9(2):41–76, 1975.

[GM76] D. Gabay and B. Mercier. A dual algorithm for the solution of nonlinear variational problems via finite element approximation. *Computers & Mathematics with Applications*, 2(1):17–40, 1976.

[GO09] T. Goldstein and S. Osher. The split Bregman method for $L1$-regularized problems. *SIAM J. Imaging Sci.*, 2(2):323–343, 2009.

[HUL93] J.-B. Hiriart-Urruty and C. Lemaréchal. *Convex Analysis and Minimization Algorithms. II*, volume 306 of *Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]*. Springer-Verlag, Berlin, 1993.

[HZLM11] J. Huang, S. Zhang, H. Li, and D. Metaxas. Composite splitting algorithms for convex optimization. *Computer Vision and Image Understanding*, 115(12):1610–1622, 2011.

[JMOB11] R. Jenatton, J. Mairal, G. Obozinski, and F. Bach. Proximal methods for hierarchical sparse coding. *Journal of Machine Learning Research*, 12(1):2297–2334, 2011.

[Kiw85] K. C. Kiwiel. *Methods of Descent for Nondifferentiable Optimization*, volume 1133 of *Lecture Notes in Mathematics*. Springer-Verlag, Berlin, 1985.

[KRR99] K. C. Kiwiel, C. H. Rosa, and A. Ruszczynski. Proximal decomposition via alternating linearization. *SIAM Journal on Optimization*, 9(3):668–689, 1999.

[LM79] P.-L. Lions and B. Mercier. Splitting algorithms for the sum of two nonlinear operators. *SIAM J. Numer. Anal.*, 16(6):964–979, 1979.

[LPR14] X. Lin, M. Pham, and A. Ruszczyński. Alternating linearization for structured regularization problems. *The Journal of Machine Learning Research*, 15(1):3447–3481, 2014.

[PR55] D. W. Peaceman and H. H. Rachford. The numerical solution of parabolic and elliptic differential equations. *J. Soc. Indust. Appl. Math.*, 3:28–41, 1955.

[Roc76] R. T. Rockafellar. Monotone operators and the proximal point algorithm. *SIAM Journal on Control and Optimization*, 14(5):877–898, 1976.

[Rus06] A. Ruszczyński. *Nonlinear optimization*, volume 13. Princeton University Press, 2006.

[RW91] R. T. Rockafellar and R. J.-B. Wets. Scenarios and policy aggregation in optimization under uncertainty. *Mathematics of operations research*, 16(1):119–147, 1991.

[VRMV14] S. Villa, L. Rosasco, S. Mosci, and A. Verri. Proximal methods for the latent group lasso penalty. *Computational Optimization and Applications*, 58(2):381–401, 2014.

[WBL14] H. Wang, A. Banerjee, and Z.-Q. Luo. Parallel direction method of multipliers. pages 13–64. Neural Information Processing System (NIPS), 2014.

[Yu13] Y.-L. Yu. Better approximation and faster algorithm using the proximal average. In *Advances in Neural Information Processing Systems*, pages 458–466, 2013.

[YX11] G.-B. Ye and X. Xie. Split Bregman method for large scale fused Lasso. *Comput. Statist. Data Anal.*, 55(4):1552–1569, 2011.